

Listing of All Claims

1. (Currently amended) A method for combining extents of a source data set, with minimal down time imposed on applications accessing the source data set, the source data set having a corresponding name, source DASD device(s) and one or more extent locations on the source DASD device(s), the method comprising the steps of:
- a) assessing the source data set to be migrated, the volume(s) it resides on, and the total space it occupies;
 - b) allocating space for a corresponding target data set using the total source space as a primary allocation request and using the original source data set name;
 - c) designating one or more target locations in the allocated space and assigning each source data set extent location to a respective one or more of the designated target locations;
 - d) starting monitor programs on any images that can write to any of the identified source volumes in order to detect a subsequent change to the source data set;
 - e) storing an indication of each source data set track detected by a monitor program as having changed;
 - f) begin copying the source data set in accordance with said assignment of each source data set extent so as to form the corresponding target data set;
 - g) while copying the source data set, periodically re-synchronizing source and target tracks detected by the monitor programs as having changed after having been previously copied;
 - h) upon substantial completion of said copying, signaling ~~the said application programs~~ an operator that closing the source data set is requested;
 - i) recognizing that the source data set is closed thereby commencing a down time window;
 - j) during the downtime window, finally re-synchronizing the source and target tracks detected by the monitor programs as having changed after having been previously copied;

- k) during the downtime window, accommodating any allocation ~~differences~~ changes affecting the source data set that occurred since the initial assessment of the source data set;
- l) changing catalog entries to reflect new target data set volumes; and
- m) signaling the ~~application(s)~~ user that the target data set may be opened, thereby terminating the down time window.

2. (original) A method according to claim 1 wherein said designating one or more target locations in the allocated space and assigning each source data set extent to a respective target location includes constructing a cylinder/track translate table having an entry for each extent of the source data set.

3. (original) A method according to claim 1 wherein said accommodating allocation differences includes detecting and accommodating a persistent data set.

4. (original) A method according to claim 1 wherein said accommodating allocation differences includes detecting and accommodating a new source data set added during the mirroring process.

5. (original) A method according to claim 1 wherein said accommodating allocation differences includes detecting and accommodating a source data set deleted during the mirroring process.

6. (original) A method according to claim 1 wherein said accommodating allocation differences includes detecting and accommodating extent changes to data sets included in the mirroring process.

7. (new) A method for data set-level mirroring of a selected source data set while minimizing downtime of applications accessing the source data set, the method comprising:

identifying extents allocated to the source data set, each extent comprising a set of at least one DASD track;

allocating within a target data set at least a number of extents needed to store the identified extents of the source data set;
copying each track of an extent of the source data set to an extent of the target data set;
in response to any tracks of the source data set being changed by an application during the copying step, recopying the changed tracks from the source data set to the target data set;
commencing a downtime window in which applications may no longer access the source data set; and
in response to an allocation change being made to the source data set since the allocating step, performing a corresponding allocation change to the target data set during the down-time window.

8. (new) The method of claim 7, wherein allocating comprises assigning each extent of the source data set to a respective extent of the target data set.

9. (new) The method of claim 8, wherein assigning comprises constructing a cylinder/track translate table having an entry for each extent of the source data set that points to the respective extent of the target data set.

10. (new) The method of claim 9, further comprising flagging an entry of the cylinder/track translate table when a track of a corresponding extent of the source data set has been changed.

11. (new) The method of claim 10, wherein recopying the changed tracks comprises recopying each track having a flagged entry in the cylinder/track translate table.

12. (new) The method of claim 7, wherein the source data set has a name, and wherein identifying comprises:
identifying within a catalog a data set having a matching name; and
determining from the catalog one or more volumes used by the identified data set.

13. (new) The method of claim 12, wherein identifying extents comprises gathering extent information from at least one Volume Table of Contents (VTOC) for the one or more volumes.
14. (new) The method of claim 12, wherein allocating comprises allocating the extents for the target data set without cataloging such that the target data set and the source data set may have the same name.
15. (new) The method of claim 7, wherein allocating comprises storing meta-data relating to the source data set, and wherein performing an allocation change comprises determining that an allocation change has been made to the source data set by comparing the stored meta-data with current meta-data for the source data set.
16. (new) The method of claim 15, wherein the meta-data comprises at least one of a catalog, a Volume Table of Contents (VTOC), and a VVDS (VSAM Volume Data Set).
17. (new) The method of claim 7, wherein performing an allocation change comprises allocating one or more additional extents within the target data set for any new extents added to the source data set after the allocating step.
18. (new) The method of claim 17, further comprising copying each track of a new extent of the source data set to one of the additional allocated extents of the target data set.
19. (new) The method of claim 7, wherein performing an allocation change comprises deleting extents within the target data set in response to a corresponding extent within the source data set being deleted after the allocating step.
20. (new) The method of claim 19, wherein deleting extents comprises freeing volume space for deleted extents.

21. (new) The method of claim 7, wherein performing an allocation change comprises renaming the target data set in response to the source data set being renamed after the allocating step.

22. (new) The method of claim 7, wherein commencing a downtime window comprises closing the source data set.

23. (new) The method of claim 7, wherein commencing a downtime window comprises notifying one or more application programs that the source data set is to be closed.

24. (new) The method of claim 7, further comprising initiating a process for monitoring for changes to the tracks prior to the step of copying the tracks.

25. (new) The method of claim 7, wherein recopying comprises periodically recopying changed tracks from the source data set to the target data set.

26. (new) The method of claim 7, further comprising recopying any changed tracks from the source data set to the target data set after the downtime window has commenced.

27. (new) An article of manufacture comprising a computer-readable medium storing computer-readable program code for performing a method for data set-level mirroring of a source data set while minimizing downtime of applications accessing the source data set, the article of manufacture comprising:

- computer-readable program code for identifying extents belonging to the source data set, each extent comprising a set of tracks;
- computer-readable program code for allocating within a target data set at least a number of extents needed to store the identified extents of the source data set;
- computer-readable program code for copying each track of an extent of the source data set to an extent of the target data set;

computer-readable program code for in response to any tracks of the source data set being changed by an application during the copying step, recopying the changed tracks from the source data set to the target data set;

computer-readable program code for commencing a downtime window in which applications may no longer access the source data set; and

computer-readable program code for performing an allocation change to the target data set in response to an allocation change being made to the source data set since the initial allocating step.

28. (new) The article of manufacture of claim 27, wherein the computer-readable program code for allocating comprises program code for assigning each extent of the source data set to a respective extent of the target data set.

29. (new) The article of manufacture of claim 28, wherein the computer-readable program code for assigning comprises computer-readable program code for constructing a cylinder/track translate table having an entry for each extent of the source data set that points to the respective extent of the target data set.

30. (new) The article of manufacture of claim 29, further comprising computer-readable program code for flagging an entry of the cylinder/track translate table when a track of a corresponding extent of the source data set has been changed.

31. (new) The article of manufacture of claim 30, wherein the computer-readable program code for recopying the changed tracks comprise computer-readable program code for recopying each track having a flagged entry in the cylinder/track translate table prior to commencing a downtime window.

32. (new) The article of manufacture of claim 27, wherein the source data set has a name, and wherein the computer-readable program code for identifying comprise:

computer-readable program code for identifying within a catalog a data set having a matching name; and

computer-readable program code for determining from the catalog one or more volumes used by the identified data set.

33. (new) The article of manufacture of claim 32, wherein the computer-readable program code for identifying extents comprise computer-readable program code for gathering extent information from at least one Volume Table of Contents (VTOC) for the one or more volumes.

34. (new) The article of manufacture of claim 32, wherein the computer-readable program code for allocating comprise computer-readable program code for allocating the extents for the target data set without cataloging such that the target data set and the source data set may have the same name.

35. (new) The article of manufacture of claim 27, wherein the computer-readable program code for allocating comprise computer-readable program code for storing meta-data relating to the source data set, and wherein the computer-readable program code for performing an allocation change comprise computer-readable program code for determining that an allocation change has been made to the source data set by comparing the stored meta-data with current meta-data for the source data set.

36. (new) The article of manufacture of claim 35, wherein the meta-data comprise at least one of a catalog, a Volume Table of Contents (VTOC), and a VVDS (VSAM Volume Data Set).

37. (new) The article of manufacture of claim 27, wherein the computer-readable program code for performing an allocation change comprise computer-readable program code for allocating one or more additional extents within the target data set for any new extents added to the source data set after the allocating step.

38. (new) The article of manufacture of claim 37, further comprising computer-readable program code for copying each track of a new extent of the source data set to one of the additional allocated extents of the target data set.

39. (new) The article of manufacture of claim 27, wherein the computer-readable program code for performing an allocation change comprise computer-readable program code for deleting extents within the target data set in response to a corresponding extent within the source data set being deleted after the allocating step.

40. (new) The article of manufacture of claim 39, wherein the computer-readable program code for deleting extents comprise computer-readable program code for freeing volume space for deleted extents.

41. (new) The article of manufacture of claim 27, wherein the computer-readable program code for performing an allocation change comprise computer-readable program code for renaming the target data set in response to the source data set being renamed after the allocating step.

42. (new) The article of manufacture of claim 27, wherein the computer-readable program code for commencing a downtime window comprise computer-readable program code to first confirm closure of the source data set.

43. (new) The article of manufacture of claim 27, wherein the computer-readable program code for commencing a downtime window comprise computer-readable program code for requesting that the source data set be closed.

44. (new) The article of manufacture of claim 27, further comprising computer-readable program code for initiating a process for monitoring for changes to the tracks prior to the step of copying the tracks.

45. (new) The article of manufacture of claim 27, wherein the computer-readable program code for recopying comprise computer-readable program code for periodically recopying changed tracks from the source data set to the target data set.

46. (new) The article of manufacture of claim 27, further comprising computer-readable program code for recopying any changed tracks from the source data set to the target data set after the downtime window has commenced.

47. (new) A method for data set-level migrating a selected source data set under an MVS operating system while minimizing downtime of the source data set, the method comprising:

- (a) assessing the source data set to be migrated so as to identify a set of source data set tracks on which the data set is stored;
- (b) storing initial meta-data that reflects extent locations of the source data set tracks;
- (c) initially allocating a target data set with at least sufficient space to store the source data set tracks;
- (d) beginning sequential copying of only the source data set tracks to the target data set;
- (e) while the sequential copying proceeds, beginning monitoring to detect any changes to a source data set track after the track was last sequentially copied, thereby identifying as a modified track one that requires re-copying;
- (f) continuing the sequential copying of source data set tracks until all of the source data set tracks have been copied, and all modified tracks have been re-copied;
- (g) after completion of the sequential copying, requesting that the user close and unallocate the source data set;
- (h) continuing said monitoring to detect and identify any modified source data set tracks until the source data set has been closed and unallocated;
- (i) while waiting for the source data set to be closed and unallocated, re-copying any identified modified track;
- (j) responsive to an indication that the source data set has been closed and unallocated, effecting a final re-synchronization step by re-copying any modified track not yet re-copied;
- (k) checking to detect any allocation change to the source data set since initially allocating the target data set, by comparing current meta-data describing the source data set to the initial meta-data stored at the beginning of the process;

(l) in response to detecting an allocation change having been made to the source data set after the target data set was initially allocated, adjusting the target data set to accommodate the allocation change;

(m) updating a corresponding catalog and system meta-data describing the source data set to point to the target data set; and then

(n) signaling release to open the target data set.

48. (new) A method for data set-level migrating according to claim 47 wherein the detected allocation change to the source data set after initially allocating the target data set comprises an extension to the source data set, and said adjusting the target data set comprises, if necessary, extending the target data set so as to accommodate the source data set extension.

49. (new) A method for data set-level migrating according to claim 47 wherein the detected allocation change to the source data set after initially allocating the target data set comprises a release of space in the source data set, and said adjusting the target data set comprises releasing corresponding space in the target data set.

50. (new) A method for data set-level migrating according to claim 47 wherein the detected allocation change to the source data set after initially allocating the target data set indicates that the source data set was removed, and said adjusting the target data set comprises releasing all space initially allocated to the target data set.

51. (new) A method for data set-level migrating according to claim 47 wherein the detected allocation change comprises allocation of a new source data set, and said adjusting the target data set comprises allocating target space for the new source data set and then copying the tracks comprising the new source data set to the corresponding target space.

52. (new) A method for data set-level migrating according to claim 47 wherein said storing initial meta-data that reflects extent locations of the source data set tracks comprises storing a cylinder/track translate table.

53. (new) A method for data set-level migrating according to claim 47 wherein said initially allocating a target data set includes using the source data set name as the allocated target data set name.

54. (new) A method for data set-level migrating according to claim 47 wherein the said checking step to detect any allocation change to the source data set is effected after the final re-synchronization step.

55. (new) A method for data set-level migrating according to claim 47 and further comprising effecting a preliminary re-synchronization by re-copying the modified tracks prior said requesting that the user close and unallocate the source data set.